

Using Model-Based Reasoning for Autonomous Instrument Operation¹

Michael A. Johnson
NASA/Goddard Space Flight Center
Code 564
Greenbelt, MD 20771
301 286 3170
michael.a.johnson.2@gsfc.nasa.gov

Michael L. Rilee
NASA/ Goddard Space Flight Center -RAYTHEON
Code 930
Greenbelt, MD 20771
301 286-4743
michael.rilee@gsfc.nasa.gov

Walt Truskowski
NASA/Goddard Space Flight Center
Code 588
Greenbelt, MD
301 286-8821
Walt.Truskowski@gsfc.nasa.gov

Abstract—Model-based reasoning has been applied as an autonomous control strategy on the Low Energy Neutral Atom (LENA) instrument currently flying on board the Imager for Magnetosphere-to-Aurora Global Exploration (IMAGE) spacecraft. Explicit models of instrument subsystem responses have been constructed and are used to dynamically adapt the instrument to the spacecraft's environment. These functions are cast as part of a virtual-Principal Investigator (VPI) that will, when complete, autonomously monitor and control the instrument. Even in the VPI's current partial implementation, LENA's command uplink volume has been decreased significantly from its previous volume. This work demonstrates that a model-based approach can be used to enhance science instrument effectiveness. The components of LENA are common in space science instrumentation, and lessons learned by modeling this system may be applied to other instruments. Future work involves the extension of these methods to cover more aspects of LENA operation and the generalization to other space science instrumentation.

TABLE OF CONTENTS

1. INTRODUCTION
2. AUTONOMY OPTIONS
3. MODEL-BASED REASONING APPLIED TO IMAGE/LENA
4. MODEL-BASED REASONING ISSUES AND CHALLENGES
5. CONCLUSION

I. INTRODUCTION

Multiprobe missions are an important part of NASA's future. Consider the missions of the Sun-Earth Connections (SEC) theme, which include missions such as Magnetospheric Multi Scale (MMS, five spacecraft, launch 2006) and the Magnetospheric Constellation Draco (50-100 spacecraft, launch 2010). Members of NASA's Solar Terrestrial Probe line, these missions are part of a series of technologically ambitious projects that build towards the placement of a distributed sensor web that can accurately measure the mesoscale structure and dynamics of Geospace. Geospace is the region of space wherein the Sun and Earth interact to produce Space Weather. To make such missions robust, reliable, and affordable, ideally the many spacecraft of a constellation must be at least as easy to operate as one spacecraft is today.

This level of performance is to be achieved in spite of full suites of scientific instruments, limited communication opportunities perhaps separated by weeks, and limited ground operations resources. Downlink bandwidth limitations reduce the coverage and resolution of the science products that missions may produce. Furthermore, understanding many important phenomena requires simultaneous measurements from multiple spacecraft. Operations techniques that require communication with the ground incur communications latencies and suffer bandwidth limitations that inhibit a mission's ability to react

¹ U.S. Government work not protected by U.S. copyright.

to science of opportunity, to coordinate collective behaviors across the constellation, and to deal with faults.

The advantages of spacecraft autonomy have been perceived for some time, and for early missions such as Ranger 6[1] and the early Mars and Venus Mariners[2] an amount of autonomy was a matter of course due to low communication rates and limited commandability. Advances in space borne computers and communications technology led to spacecraft that could more readily be configured and commanded. Part of this trend has continued as computer technology has presented opportunities first to automate, and then to add flexibility and fault tolerance to different segments of the mission[3]. Increasing numbers of increasingly complex spacecraft have led to the recent study and application of even more sophisticated approaches [ref. RAX, ref. Shuttle, ref. HST].

One recent approach that is relevant to science missions is the Remote Agent Executive (RAX) experiment that operated Deep Space One with some success during an asteroid flyby[4]. One module of RAX maintains a set of models that correspond to spacecraft systems and makes plans and resolves conflicts by reasoning based on these models[5]. A key driver of RAX resource requirements is the complexity of constructing mission plans that maintain system constraints. One way to reduce this complexity is to delegate responsibility for operation to spacecraft subsystems, leading to the concept of subsystem or instrument-based autonomy[6].

To reduce the complexity of the problem we feel it best to aggressively attack and reduce system complexity at each level of a system's hierarchy. Reducing system complexity at the lowest levels may dramatically reduce the complexity of the overlying control functions. This eases the burden of spacecraft level autonomy, e.g. at the level of a spacecraft agent like RAX. By moving instrument operations as far to the instrument as possible, including the autonomous production of data products, communication resource requirements can be dramatically reduced while at the same time dramatically improving the quality and quantity of the science obtained[7].

We have tested these ideas in the context of the Low Energy Neutral Atoms (LENA) experiment that is flying on the Imager for Magnetosphere-to-Aurora Global Exploration (IMAGE) observatory[8]. IMAGE is a NASA/SEC mission designed to obtain a global picture of the Earth's magnetosphere using a variety of remote sensing techniques. LENA, being a particle detector, can be impaired or may fail because of excessive particle fluxes or environmental radiation[9]. We have constructed an explicit model of LENA's response. The instrument uses this model to dynamically adapt its response to autonomously maintain instrument health and safety and improve science return. We call the reasoning system that uses the model to determine how to configure LENA the Virtual Principal-Investigator (VPI) because of the responsibility it holds for

instrument operations. By implementing these functions at the instrument level, it was possible to bring these advanced behaviors into the very constrained computing environment of the LENA Flight Model. Furthermore, these enhancements were realized with no deleterious impact on other IMAGE systems.

In this paper we are focusing on a proposed approach to achieve autonomy for scientific instruments. We begin by considering some autonomy options and by presenting an overview of the model-based approach to autonomous instrument operations that is currently under investigation. Then we discuss the application of these ideas to LENA followed by a discussion of challenges to generalizing our model-based approach to other instruments. We close with a discussion of future work along these lines.

2. AUTONOMY OPTIONS

NASA mission systems have two major components: ground-based command/control and science operations; and space-based flight-software and instrument systems. Each of these components has complex real-time operational constraints that must be met in order to ensure mission safety and science agenda success. Much has been done to automate each of these components.

Ground-based Command/Control and Science Operations

Much attention has been paid to realizing ground-system autonomy. Through the use of expert-system technology and scripting languages, a nearly complete lights-out approach to ground-based operations has been achieved for certain NASA missions. In these systems human intervention is only required in certain extreme situations that the automated (operator-attended) system is unable to handle. Additionally, Science Operations centers are usually equipped with advanced tools to enable rapid science planning and science agenda management.

Space-based Control

The state of onboard autonomy is also relatively high for some onboard functions. Over the years much autonomy has been realized onboard through the evolution of the flight-software system. Table 1 gives an indication of the autonomy levels currently available for typical onboard activities[10].

Table 1 – Current Onboard Autonomy	
Onboard Activity	Current Level of Autonomy
Planning & Scheduling	low
Command Loading	n/a
Science Schedule Execution	medium
Science Support Activity Execution	medium
Onboard Engineering Support Activities	high
Downlinked Data Capture	n/a
Data and Performance Monitoring	medium
Fault Diagnosis	low
Fault Correction	low
Downlinked Data Archiving	n/a
Engineering Data Analysis/Calibration	low
Science Data Processing/Calibration	low

Model-based Autonomous Instrument Operations

The focus of this paper is on future autonomy for spacecraft instrument operations. A current study, reported on in this paper, focuses on a model-based reasoning approach to instrument autonomy and its application to the autonomy of the LENA instrument on the IMAGE spacecraft. Figure 1 depicts the major concepts associated with this study. The basic idea is quite straightforward. The ground-based Principal Investigator (PI - we use PI both in a singular and collective sense) has a mental model on the instrument and its designed behaviors, inputs, outputs. This model is

formalized and codified and becomes the model that an onboard intelligent process uses to guide the operations of the instrument and to aid in diagnosing instrument faults and taking corrective actions.

From our perspective and in our LENA-related work a model is some representation of reality that is used to support an understanding of that reality and to make decisions about the behaviors associated with that reality. As Kaposi and Meyers [11] put it: “A good model is not an arbitrary representation, but a systematic, purposeful simplification of its referent. Such a model focuses attention on selected attributes which hold the key to the solution of the given problem and suppresses features which, for the given situation, are less crucial or irrelevant.”

For our purposes we are dealing with models of components of the LENA instrument and their integration into an overall LENA model.

There are several ways to deal with models. Three major classifications of use are:

- Internal representation (model embedded in code as a procedure). In this case the model is implicit thus difficult to readily modify or adapt.
- External representation (model expressed in an external knowledge representation external to and processed by some procedural code). In this case the model is explicit and thus easily modified or adapted to a changing environment.
- Hybrid representation (a combination of the two).

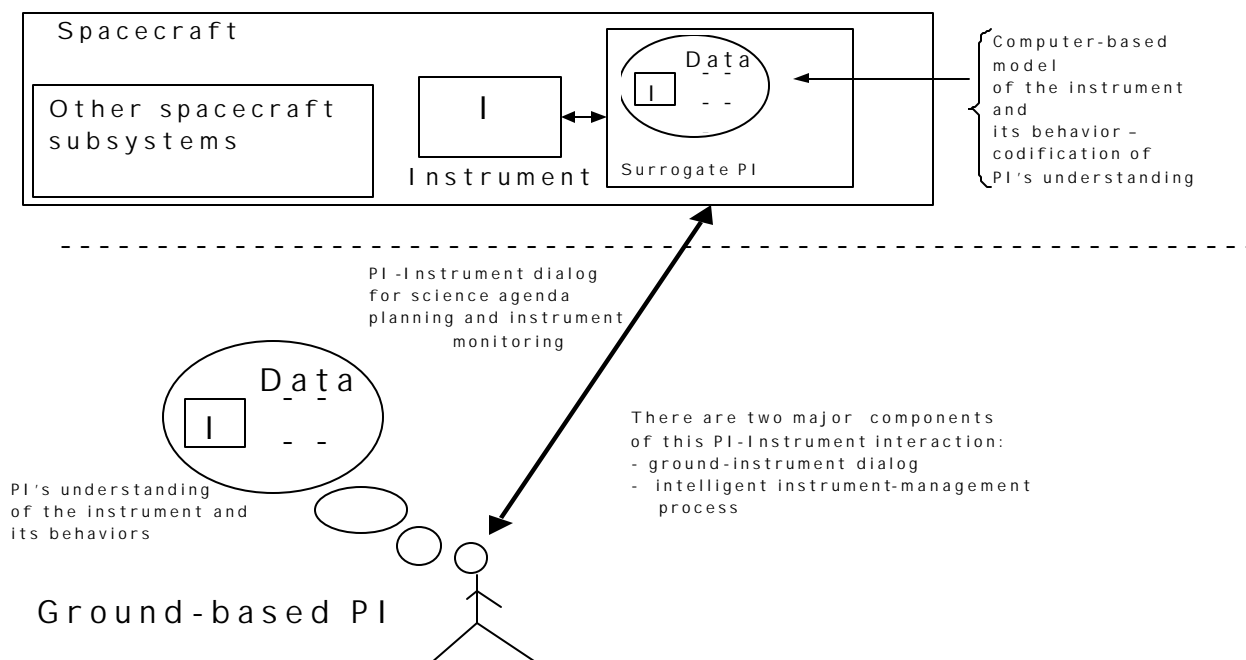


Figure 1 – Major Autonomy Concepts

In this phase of the LENA modeling work, use has been very successfully made of the internal representation approach resulting in on-board software automating various LENA functions. The intent is to graduate to the external representation approach. The following section will discuss what we have achieved so far in applying this autonomy concept to LENA.

3. MODEL-BASED REASONING APPLIED TO IMAGE/LENA

IMAGE/LENA Objectives

The IMAGE observatory is a spin-stabilized spacecraft that was launched in March 2000 into an elliptical polar orbit with an apogee altitude of 7.2 earth radii (45,922 km) and a perigee altitude of 1000 km. It is the first satellite dedicated to imaging earth's magnetosphere. LENA, one among the suite of 6 instruments on the payload uses high-voltage electrostatic optics and time-of-flight mass spectroscopy to image fast neutral atom flux and measure its composition and energy distribution.

System Implementation

Simulated particle trajectories are plotted in Figure 2.

Neutral particles (1) enter the instrument through a collimator (2) which filters charged particles. The tungsten surface (3) converts neutrals to negative. Negative ions from the surface are then collected by the extraction lens (4), which focuses all negative ions with the same energy to a fixed location. The ions are then accelerated by a high voltage optics potential prior to entering the electrostatic analyzer (5). Finally, the ions pass into a time-of-flight/position sensing section (6) where ion mass, energy, and angle are determined[12].

The electrostatic potentials required to conduct the experiment are derived from 5 commandable high-voltage power supplies- 2 collimator supplies, an ion optics supply and 2 microchannel plate (MCP) supplies. These supplies and the TOF subsystem are controlled by an 8-bit microcontroller-based command and data handling system.

IMAGE is operated as a lights-out mission- the spacecraft is out of ground contact except for once during each 14.2-hour orbit. To support this operations paradigm, the IMAGE central instrument data processor permits queued commands to be routed to LENA at predetermined times. This allows the instrument to be configured in response to predictable conditions. It is important however, that LENA also have the capability to react immediately to conditions that could threaten the health and safety of instrument systems- autonomous real-time command capability.




Figure 2 - Neutral Atom Ray-Tracing

Autonomous Operations Virtual Principal Investigator (VPI)

The ground-based PI's ability to configure the instrument in response to dynamic conditions is hindered by limited observability of LENA parameters. The downlink bandwidth allocated to LENA renders it unfeasible to telemeter parameters at a sample-rate high enough to ensure that transient behavior will be captured. Communications latency further constrains real-time responses. We address these issues by conveying a subset of LENA's command authority from the ground to a Virtual Principal Investigator (VPI) onboard the instrument (Figure 3).

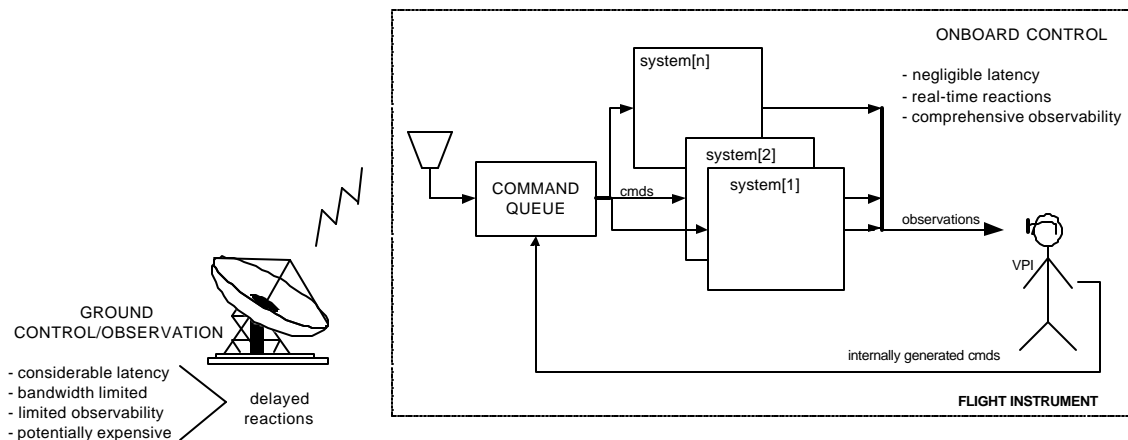


Figure 3 - Onboard vs. Ground-Based Control

The VPI provides the capability to respond in real-time to predicted (e.g. radiation belt) and random (e.g. solar storms) conditions. Actions that can be initiated onboard are consistent with the command authority granted by the ground-based PI (Table 2).

Table 2 – Onboard Authority Level		
Authority Level Granted	Reasoning Locale	Possible Onboard Actions
I	onboard	Initiate commands Inform ground of actions taken
II	Onboard/ ground	Submit recommended actions to ground
III	ground	none

The VPI is primarily tasked with monitoring and controlling two critical LENA behaviors: instrument over-stimulation and high-voltage health and safety. Potentially damaging event rates could result from high-flux environments. They could also be indicative of high-voltage discharges that could degrade electrostatic surfaces and damage electronic components. In either case, the start or stop channel gains must be reduced to limit the resultant count rates. Operation of the high-voltage systems is also monitored. The status of each high-voltage supply is thereby derived. Furthermore, the state of the electrostatic surfaces can be indirectly inferred since excessive currents or unregulated voltages may be indicative of anomalous conditions on these surfaces. Control of these behaviors is granted authority level I.

Model-based Reasoning

The VPI operates within a model-based framework. The behavior of the instrument is decomposed into a family of behavioral models with each model mapped to a subsystem. A model captures the electrical response function of the targeted system. The models are typically excited with the same stimuli as the systems they represent. The resultant responses are routed to the VPI, which considers whether the current instrument state is desirable, and if not, initiates corrective actions.

High Voltage Power Subsystem (HVPS)—The HVPS models incorporate components of varying complexity. The degree of complexity is consistent with information the VPI requires to ascertain and control the state of system. A first-order polynomial appropriately codifies the voltage response of the HVPS to commands whereas the current response is represented as a constant source; the current-threshold test implemented by the VPI does not require a high degree of model fidelity (Figure 4).

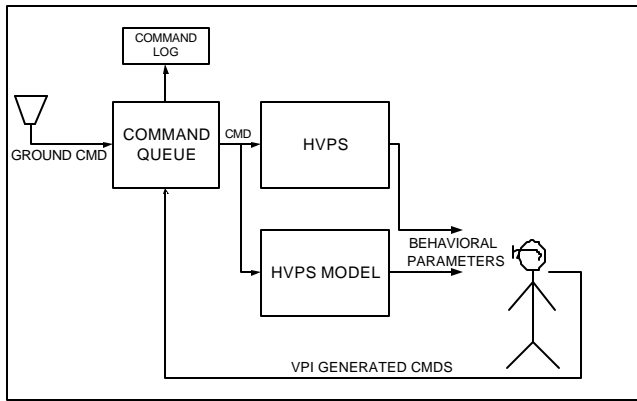


Figure 4 - Virtual-PI HVPS Control

The VPI uses the model outputs to maintain or correct the state of LENA. Measured and modeled parameters are updated and reaction commands are executed on a time-scale consistent with the dynamics of the targeted system and the control objectives of the VPI. For example, hvps reasoning is implemented as follows: compare the voltage response of the power supply with the expected response. If the deviation exceeds P_1 volts for longer than P_2 seconds, take P_3 action. Here, P_n are parameters that can be varied under ground control. Furthermore, compare measured currents and voltages to the threshold levels P_4 . If a threshold is exceeded for longer than P_5 seconds, take action P_6 , where again P_n are ground commandable parameters.

Time of Flight (TOF) Subsystem—Overstimulation of the start or stop MCP channels can compromise the TOF system. Excessive rates could result from the periodic radiation passes or from energetic solar ions.

The effective gains of the TOF channels are proportional to the MCP-start and MCP-stop potentials. While science return is maximized when the MCP voltages are at their nominal levels, high count rates are also most likely to occur. The goal of the VPI is therefore to maximize nominal-level operations, but reduce voltages as required to maintain instrument health and safety.

Time-tagged commands could be used to decrease the gains during the periodic radiation passes. The drawback of this approach however is that count rates cannot be predicted as a function of time with great accuracy. Therefore, fairly conservative time boundaries are typically used to define when instruments should safe themselves. This approach compromises science return.

A more robust approach is to react directly to count rates. Gains are reduced only when required. This approach has the advantage of not only reacting to events that result from the periodic radiation encounters, but to unpredictable energetic particle events as well.

After the VPI has configured the instrument to protect itself in response to a high-rate scenario, it must determine when normal operations can resume. Since the operational voltages have been reduced, measured count rates cannot be used directly in this determination. Instead, a model of each channel is used to predict when the voltages can be increased to nominal levels without violating an overstimulation criteria.

An occurrence when the flight system was overstimulated is shown in Figure 5.




Figure 5 – Stop MCP Channel Overstimulation Response

The stop singles response shown in the upper plot is a function of the incoming particle flux and the corresponding stop MCP voltage shown in the lower plot. The VPI detected instrument overstimulation and subsequently reduced the voltage from the nominal operating point as shown in the lower plot.

An MCP gain model is used to predict when the nominal operating voltage can be restored (Figure 6).




Figure 6 – MCP Gain Model $M(\tau, v)$

The model, codified as a 2-dimensional lookup table is parameterized with respect to MCP HVPS voltage and a signal detection threshold τ used in the TOF electronics. The model acts upon the measured flux to predict the count

rate R that would be measured at the nominal operating voltage v_{nom}

$$R(\text{predicted}) = k \cdot \text{flux}_{\text{measured}} \cdot M(\tau, v_{nom})$$

where k incorporates various scale factors resulting from the electronics signal processing path. The VPI restores nominal operations when the predicted rate does not exceed the overstimulation criteria.

VPI Adaptability - The set points used by the VPI to implement reasoning are normally static, although they can be changed via ground command. Optimally however, they should be adapted as a function of long-term observed behavior. The VPI facilitates this by gathering long-term operational statistics. These statistics are used to recommend optimal set points to the ground-PI (Authority level II).

4. MODEL-BASED REASONING ISSUES AND CHALLENGES

Section 3 above has detailed the status of the work on the LENA VPI and discussed the model-based approach toward realizing its incremental autonomy. In this section we list some abstract issues associated with model-based reasoning that were influenced by our initial and continuing attempts to realize a true model-based approach to instrument autonomous operations. Addressing these issues will significantly help us in our ongoing project efforts.

The modeling issues that are identified in this section are certainly not new to computer science. Much of the current model-based reasoning literature addresses them and currently, efforts are underway to identify the most appropriate approaches for handling these issues in our context. The recent work on model-based reasoning for the Remote Agent Project at the Ames Research Center (ARC) [13] and the work of Bernard Zeigler et. al. [14] on modeling and simulation are but two of the many resources that are being investigated.

From the standpoint of model-based reasoning as a technology there are many challenging issues to be addressed[15]. We briefly discuss three:

- Adaptability
- Granularity
- System scoping

Adaptability

Some systems can vary over time. This is especially true for agent-based autonomous systems. Intelligent agents evolve over time based on many factors, including being embedded in and having to adapt to a changing environment. Adaptation can also occur because of an agent's ability to


learn through self assessment against performance criteria established for it. Formally stated:

Question 1. Given that $M(S_t) = T$ (i.e., the model of system S at time t is a correct model) what can formal modeling do to allow one to evolve the M to ensure that $M(S_{t+1}) = T$ where S has changed over time because of system environmental changes or "learning"?

Is there a tried and true way to evolve the model M to reflect the changes the system S has undergone because of adaptation to environmental changes or the system's learning process?

Granularity

The next question is related to what we call granularity. As systems get "more complicated" a single formal method or model may be inappropriate or not feasible for the system as a whole. System decomposition into components and the individual modeling of each component by an appropriate formal technique tailored to that component may be required. Multiple differing formal and/or heuristic models may be required to address the various nuances of the system's components. The question is: will an algebra of



models be available to allow composition of the individual component models into a model of the entire system, i.e., will we have

$$M(S) = M_1(S_1) + M_2(S_2) + \dots + M_n(S_n)$$

System Scoping

In classical cybernetic theories, a system has to have a definite boundary that delimits it from its environment. The

Figure 7 – System Scoping

question, which we raise regarding system scoping, has to do with what is included within the boundary. Figure 7 illustrates one instance of the question regarding scope.

The left hand side illustrates a person interacting with a system S according to his mental model of S . The right hand shows the user and his mental model as being an integral part of a larger system S_1 . Formal methods today address the modeling of system S . Will formal methods be able to address the modeling of system S_1 , i.e., systems in which the human is an integral component?

All three of these issues will bear on the model-based reasoning application in realizing instrument autonomy.

From a model-based reasoning application point-of-view the major challenges are related to

- reliability,
- completeness and
- usability issues.

Reliability and Completeness

In order for confidence to be placed in the use of model-based reasoning as the foundation for realizing instrument autonomy there needs to be a level of confidence on the part of ground-based personnel that the model is correct, robust, comprehensive, adaptable, and reliable. This is a tall order. The level of detail needed to support instrument autonomy needs to be clearly understood. This is a new active area of research.

Usability

The usability issue, in our context, is somewhat unique. In addition to the generation of a model (or models) that satisfy the human need to fully document instrument structure and operational behaviors, the model development and representation needs to be usable by an autonomous process such as the VPI which will use the model to reason about system behaviors and diagnose and repair instrument anomalies. Finding a model notation that satisfies the needs and capabilities of these two distinct classes of users is a major challenge.

Future Work

There are many challenges facing us in our attempt to realize a true model-based approach to LENA autonomous operations. There are two perspectives to the autonomy question: the first is realizing autonomy from an instrument subsystem (e.g. engineering) perspective; the second is to realize autonomy from the science perspective (a much more difficult task). We will continue to focus on the first perspective in the immediate future. For our future work we have two major goals:

- complete the model-based approach to LENA autonomy focusing on the instrument subsystem perspective
- generalize the approach to make it applicable to other instruments in a user-friendly manner

The success we have experienced so far puts us in a very good position to achieve both these goals. Time will tell.

ACKNOWLEDGEMENTS

The authors would like to thank the IMAGE/LENA Principal Investigator, Dr. Thomas Moore of NASA/Goddard Space Flight Center for his support of this effort. We'd also like to thank Dr. Dennis J. Chornay also from GSFC for his assistance in characterizing the TOF system.

REFERENCES

- [1] W. E. Burrows, *Exploring Space: Voyages in the Solar System and Beyond*. Random House, N.Y., p 78, 1990.
- [2] M. W. Hunter and R. L. Moll, "Future Interplanetary Unmanned Space Missions." In *Future Space Program and Impact on Range and Network Development*, v. 15, of AAS Science and Technology Series, Morganthaler, G. W. (ed) p. 114, 1967.
- [3] J. C. Horvath, "Spacecraft Autonomy Issues: Present and Future", British Interplanetary Society Journal, 1995.
- [4] B. Pell et al., "Mission Operations with an Autonomous Agent", In Proceedings of the IEEE Aerospace Conference, Snowmass, CO, 1998.
- [5] B. C. Williams and P. P. Nayak, "A Model-Based Approach to Reactive Self-Configuring Systems", Procs. Of AAAI-96, 971-978, Cambridge, Mass., AAAI, 1996.
- [6] M. A. Johnson, et al., "Nanosat Intelligent Power System Development", In Procs. of 2nd International Conference on Integrated Micro-Nanotechnology for Space Applications, E. Y. Robinson (ed.), v2 p475, 1999.
- [7] S. A. Curtis et al., "Small Satellite Constellation Autonomy via Onboard Supercomputing and Artificial Intelligence", In Proc. of the 51st Congress of the International Astronautical Federation, 2000.
- [8] J. L. Burch ed., "The IMAGE Mission", Space Sci. Rev., v91, Nos. 1-2, 2000.
- [9] T. Moore et al., "The Low-Energy Neutral Atom Imager for IMAGE", Space Sci. Rev., v91, p155-195, 2000.
- [10] Agent-based Spacecraft Autonomy Draft 1, W. Truszowski, L. Hallock, C. Rouff, J. Karlin, Internal GSFC Report (available upon request)
- [11] A. Kaposi and M. Myeres, "Systems, Models and Measures", Springer-Verlag, 1994.
- [12] T. Moore et al., "The Low-Energy Neutral Atom Imager for IMAGE", Space Sci. Rev., v91, p155-195, 2000.
- [13] <http://ase.arc.nasa.gov/jpf/papers.html>.
- [14] B. Zeigler, H. Praefer, T. Kim "Theory of Modeling and Simulation Second Edition", Academic Press, 2000.
- [15] Panel Discussion on the Future of Formal Methods in the First International Workshop on Formal Approaches to Agent-based Systems (FAABS), GSFC, April 2000, to be published by Springer-Verlag.

Michael A. Johnson is a senior electrical engineer in the GSFC Microelectronics and Signal Processing Branch. He has served as the lead electrical engineer for numerous flight systems at GSFC, including the IMAGE/LENA Electronics Systems the Cassini/CAPS Spectrum Analyzer, and the Triana/PlasMag Electron Spectrometer Controller. His current research interest involves incorporating autonomy into platforms with severely limited processing capabilities.



areas of human factors of website design/use and the application of agents for the intelligent access and management of web-based information.

He is associated with the Small Business Innovative Research (SBIR) program at Goddard and serves as the Goddard subtopic manager for the "Flight/Ground System Autonomous Operations" and "Automated Reasoning for Autonomous Systems" subtopics.

He is a National Research Council (NRC) accredited Fellow participating in the Resident Researcher's Associate (RRA) program at the Goddard Space Flight Center.

Prior to coming to Goddard, he was a lead electrical engineer at MIT/Lincoln Laboratory responsible for the design and development of space and ground based systems. He received his B.S, M.S and Electrical Engineer degrees from MIT.

Michael L. Rilee is a scientist with Raytheon supporting the GSFC Laboratory for Extraterrestrial Physics and the Science Computing Branch of the GSFC Earth and Space Science Computing Division. For NASA's Remote Exploration and Experimentation project he has lead development of the Plasma Moment Application and the Radio Astronomical Imager which are science data analysis applications designed for space borne supercomputers. He is currently researching a High Performance Computing System that may fly on Magnetospheric Multi Scale (launch 2007). At GSFC he has been active in Nano-Satellite technology development and the application of parallel computing to data analysis and astrophysical fluid simulation (PARAMESH). He received his Ph.D. and M.S. in Astrophysics (Plasma Physics) from Cornell University, and his B.A. in Astrophysics and Mathematics from the University of Virginia in Charlottesville, VA.



Walt Truszkowski is currently the Senior Technologist in the Advanced Architectures and Automation Branch at NASA's Goddard Space Flight Center. In that capacity he is responsible for managing the agent technology research for the Branch. Currently work is underway to establish an agent-based system for the ESA/NASA satellite SOHO. He also serves as the Lead of the Information Technology Research Group in the Branch. In that capacity, he is leading an effort to create a repository of information on technologies of importance to researchers in the organization. He is also leading the research in the



